

REPORT DOCUMENTATION PAGE

AFRL-SR-BL-TR-01-

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any of this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (04302). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

including the
reducing
'2202-
currently

0003

1. REPORT DATE (DD-MM-YYYY) 7-20-2000		2. REPORT TYPE Final performance report		3. DATES COVERED (From - To) From 4/1/1997 to 12/31/1999	
4. TITLE AND SUBTITLE Systems Based on Bayesian Belief Networks and Structural Equation Models for Command and Control Support				5a. CONTRACT NUMBER F49620-97-1-0225	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
				5d. PROJECT NUMBER	
6. AUTHOR(S) Marek J. Druzdzal, Ph.D.				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) AND ADDRESS(ES) Pittsburgh School of Information Sciences 135 North Bellefield Avenue Pittsburgh, PA 15260 Phone: 412-624-9432 Fax: 412-624-2788 Email: marek@sis.pitt.edu				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR/NM 801 N RANDOLPH STREET RM 732 ARLINGTON VA 22203-1977				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The performed project focused on a new paradigm of planning systems that are based on a combination of Bayesian networks and structural equation models. We focused on theoretical issues that surround combining the two in a practical planning system, developing the foundations for, and building a prototype of such system. The approach and the system built allow for efficient, yet normatively correct, treatment of various types of information, uncertainty, and utility. It is especially powerful in complex situations where the available information is heterogeneous and consists of a mixture of deterministic and uncertain relationships among discrete and continuous variables. Our main contributions are: (1) several fast state of the art stochastic sampling algorithms for approximate inference in graphical models, (2) treatment of reversible causal mechanisms for causal reasoning in graphical models, (3) a scheme for interactive construction of causal graphical models based on causal mechanisms, (4) an algorithm for learning graphical models from data, and (5) a prototype of the system, used by over 2,300 people world-wide.					
15. SUBJECT TERMS Bayesian networks, structural equation models, graphical models, uncertainty, decision making.					
16. SECURITY CLASSIFICATION OF: Unclassified			17. LIMITATION OF ABSTRACT UL	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Marek J. Druzdzal
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (include area code) 412-624-9432

DTIC QUALITY INSPECTED 3

20010109 095

Major Accomplishments

Major accomplishments of the project have been:

- (1) several fast state of the art stochastic sampling algorithms for approximate inference in graphical models,
- (2) treatment of reversible causal mechanisms for causal reasoning in graphical models,
- (3) a scheme for interactive construction of causal graphical models based on causal mechanisms,
- (4) an algorithm for learning graphical models from data, and
- (5) a prototype of the system, used by well over 2,000 people world-wide.

We briefly summarize each of these in the separate sections below.

Stochastic sampling algorithms

A system that is a combination of Bayesian networks and structural equation models needs to include algorithms that are flexible enough to work with both discrete (Bayesian networks) and continuous (structural equation models) variables. The algorithms have to accommodate arbitrary probability distributions and work with very large models. The only known classes of algorithms that will accommodate these requirements are stochastic sampling algorithms. In our work, we probed three directions: Latin hypercube sampling, quasi-Monte Carlo methods, and adaptive importance sampling.




Figure 1: Observed example convergence rate improvements in the proposed Latin hypercube sampling algorithm.

We proposed a scheme for producing Latin hypercube samples that can enhance any of the existing sampling algorithms in Bayesian networks. We tested this scheme in combination with the likelihood-weighting algorithm (Shachter & Peot, 1990; Fung & Chang, 1990) and showed that it can lead to a significant improvement in the convergence rate. While performance of sampling algorithms in general depends on the numerical properties of a network, in our experiments Latin hypercube sampling performed always better than random sampling. In some cases, we observed as much as an order of magnitude improvement in convergence rates. We introduced several practical ways of dealing with high storage requirements of Latin hypercube sample generation process and proposed a low-storage, anytime cascaded version of Latin hypercube sampling that introduces a minimal performance loss compared to the original scheme. Figure 1 shows the improvement in terms of mean squared error over existing methods obtained by our algorithm. We presented a paper describing the Latin hypercube sampling algorithm at the FLAIRS-2000 conference (Cheng & Druzdzal, 2000a). An earlier version of the paper won a school-wide 1999 Robert Korfhage award for the best paper co-authored between a student and a faculty member.




Figure 2: Observed example convergence rate improvements in the proposed quasi-Monte Carlo sampling (please note that the vertical scale is logarithmic).

Our second contribution in the area of sampling algorithms is investigation of a family of simulation approaches, known collectively as quasi-Monte Carlo methods based on deterministic low-discrepancy sequences. Quasi-Monte Carlo methods have been successfully applied to computer graphics, computational physics, financial engineering, and approximate integrals. They have proven their advantage in low-dimensionality problems. Even though some authors believe that the quasi-Monte Carlo methods are not suitable for problems of high-dimensionality, tests by Paskov and Traub (1995) and Paskov (1997) have shown that quasi-Monte Carlo methods can be very effective for high-dimensional integral problems arising in computational finance. Papageorgiou and Traub (1997) have reported similarly good performance in high-dimensional integral problems arising in computational physics, demonstrating that quasi-Monte Carlo methods can be superior to Monte Carlo sampling even when the sample sizes are much smaller. We were the first to apply quasi-Monte Carlo methods in Bayesian networks. We have shown that similarly to the findings in other domains, quasi-Monte Carlo methods work well in high-dimensionality problems, i.e., Bayesian networks with a large number of variables. We clarified several theoretical aspects of deterministic low-discrepancy sequences and solved practical issues related to applying them to belief updating in Bayesian networks. We proposed an algorithm for selecting direction numbers for Sobol sequence (Sobol, 1967). Our experimental results showed that low-discrepancy sequences (especially Sobol sequence) significantly improve the performance of simulation algorithms in Bayesian networks compared to Monte Carlo sampling algorithms. We presented a paper describing the quasi-Monte Carlo sampling in Bayesian networks at the UAI-2000 conference (Cheng & Druzdzel, 2000b).

Our final contribution is a dramatic performance improvement over the existing stochastic sampling algorithms for Bayesian networks in a new algorithm that we call Adaptive Importance Sampling for Bayesian networks (AIS-BN). The AIS-BN algorithm shows promising convergence rates even under extreme conditions and seems to outperform the existing sampling algorithms consistently. Three sources of this performance improvement are (1) two heuristics for initialization of the importance function that are based on the theoretical properties of importance sampling in finite-dimensional integrals and the structural advantages of Bayesian networks, (2) a smooth learning method for the importance function, and (3) a dynamic weighting function for combining samples from different stages of the algorithm. We also introduce the concept of oscillation degree, O_d , which expresses whether a network is dominated by the prior or the posterior probabilities and aids in choosing an importance function that leads to a better convergence. We tested the performance of the AIS-BN algorithm along with two state of the art general purpose sampling algorithms, likelihood weighting and self-importance sampling. We used in our tests three large real Bayesian network models available to the scientific community: with evidence as unlikely as 10^{-41} . While the AIS-BN algorithm always performed better than the other two algorithms, in majority of the test cases it achieved orders of magnitude improvement in precision of the results. Improvement in speed given a desired precision is even more dramatic, although we are unable to report numerical results here, as the other algorithms almost never achieved the precision reached even by the first few iterations of the AIS-BN algorithm.




Figure 3: Observed example convergence rate improvements in the proposed adaptive importance sampling algorithm for Bayesian networks (AIS-BN).

Figure 3 shows example performance comparison of the three algorithms. Figure 4 shows the performance of the AIS-BN algorithm at a finer scale. A paper describing the AIS-BN algorithm has been accepted by the prestigious *Journal of Artificial Intelligence Research* (Cheng & Druzdzal, 2000c). An earlier version of the paper won a school-wide 2000 Robert Korfhage award for the best paper co-authored between a student and a faculty member.




Figure 4: Observed example convergence rate improvements in the proposed adaptive importance sampling algorithm for Bayesian networks (AIS-BN): A close-up of the adaptive importance sampling algorithm in Figure 3).

Causal reversibility

We concentrated our theoretical work on defining the concept of a causal mechanism and understanding when a causal mechanism is reversible. The importance of reversible mechanisms is that they allow for encoding knowledge in terms of structural equations and conditional probability tables. This knowledge can be subsequently used in various models. Knowledge reuse reduces the modeling effort, which is crucial in adaptive interactive systems, such as those used in the military.

Causal manipulation theorems proposed by Spirtes *et al.* (1993) and Pearl (1995) in the context of directed probabilistic graphs, such as Bayesian networks, offer a simple and theoretically sound formalism for predicting the effect of manipulation of a system from its causal model. While the theorems are applicable to a wide variety of equilibrium causal models, they do not address the issue of reversible causal mechanisms, i.e., mechanisms that are capable of working in several directions, depending on which of their variables are manipulated exogenously. An example involving reversible causal mechanisms is the power train of a car: normally the engine moves the transmission which, in turn, moves the wheels; when the car goes down the hill, however, the driver may want to use the power train to slow down the car, i.e., let the wheels move the transmission, which then moves the engine. Some probabilistic systems can be also symmetric and reversible. For example, the noise introduced by a noisy communication channel does not usually depend on the direction of data transmission.

We investigated whether Bayesian networks are capable of representing reversible causal mechanisms. A conditional probability table in a Bayesian network can be viewed as a description of a causal mechanism involving a node and its direct predecessors. We studied the mathematical conditions on the tables that would allow reusing them when the causal mechanisms described by them are reversed. Building on the result of Druzdzel and Simon (1993), which showed that conditional probability tables in Bayesian networks can be viewed as descriptions of causal mechanisms, we study the conditions under which a conditional probability table can represent a reversible causal mechanism. Our analysis shows that conditional probability tables are capable of modeling reversible causal mechanisms but only when they fulfill the condition of *soundness*, which is equivalent to *injectivity* in equations. While this is a rather strong condition, there exist systems where our finding and the resulting framework are directly applicable. A paper describing our analysis has been accepted by the *Journal of Empirical and Theoretical Artificial Intelligence* (Druzdzel & van Leijen, 2000).

Causal discovery and causal manipulation

A major result from the project is a fundamental insight into the nature of causality with serious implications for causal modeling. Our work on the nature and reversibility of causal mechanisms has led us to understand the fundamental role that time plays in the direction of causality. To determine the causal structure of a static system given an external manipulation, it is necessary to look at a dynamic description of the system, i.e., a system of simultaneous differential equations (their exact form is not important, as long as we know which variables participate in which equations). This allows predicting the causal structure of the manipulated system, including possible reversal of the direction of some causal mechanisms. With respect to Bayesian networks, our finding suggests that reversible mechanisms can be described by several conditional probability tables, only one of which, determined by the structure of the system after external manipulation, is used by the model. Our work extends the "arc cutting" semantics proposed by Pearl (1991) to reversible mechanisms.

An especially troubling insight that results from our work is that equilibrium-state causal models discovered from data using the methods of *causal discovery* (e.g., Pearl, Spirtes *et al.*, 1993; Cooper & Herskovitz, 1991) cannot be used reliably for prediction of the effects of causal manipulation. Causal discovery, for the most part, is concerned with learning causal models in the form of directed acyclic graphs (DAGs) from equilibrium (as opposed to time series) data. Causal reasoning, by contrast, is concerned with using such causal DAGs to perform inferences. In particular, much work on causal reasoning has focused on the ability to predict the new probability distribution over a set of variables, V , given a causal graph $G=(V,E)$ and given the fact that some subset of variables $V' \subset V$ has been externally *manipulated* to some configuration. These types of *manipulation inferences* contrast with more common *diagnostic inferences*, in that the latter are essentially identical to Bayesian updating in a Bayesian network; whereas, the former may require the causal graph to be altered prior to performing probabilistic inference. Specifically, the ability to perform manipulation inferences is made possible by a critical postulate that we call the *Manipulation Postulate*. All formalisms for causal reasoning take the manipulation postulate as a fundamental starting point:

The Manipulation Postulate If $G=(V,E)$ is a causal graph and $V' \subset V$ is a subset of variables being manipulated, then the causal graph, G' , describing the manipulated system is such that $G'=(V,E')$, where $E' \subset E$ and E' differs from E by at most the set of arcs into V' .

In other words, manipulating a variable can cause some of its incoming arcs to be removed from the causal graph, but can effect no other change in the causal graph.

The Manipulation Theorem of Spirtes *et al.* (1992) proves that given the Manipulation Postulate and the Markov Condition, the probability distribution of the manipulated model can be calculated. Furthermore, the axiomatizations of causal reasoning of Galles and Pearl (1997) and of Halpern (1998) also take the Manipulation Postulate as a fundamental assumption.

The question that we posed in our is "Are these two lines of research (i.e., equilibrium causal discovery and manipulation reasoning) consistent?" Namely, what would happen if we took an equilibrium causal model (learned from data), and applied the manipulation formalisms to it? Are the resulting inferences guaranteed to be valid? We proved by explicit counterexample that such inferences are not guaranteed to be valid in the sense that conditional independencies in the manipulated model can differ from the conditional independencies in the

learned model of the manipulated system. Symbolically, if M_S is a learned causal model of system S , and if we use the \bullet operator to denote manipulation, then we show that $M_S \neq M_{\bullet S}$.

Our general strategy is as follows. We first present two extremely simple physical systems (an ideal gas trapped in a cylinder with a movable piston and a mass dangling from a damped spring), we show, based on physical laws what the "true" equilibrium causal graphs of these systems are. We further show that with an appropriate source of noise present in data taken from these systems, a constraint-based learning algorithm will learn the correct causal graphs. Finally, we show that the graph predicted by manipulation-type reasoning on these learned models will possess different conditional independence relations than the causal graph that would be learned from the true manipulated system. Furthermore, we show that under suitable manipulations, these systems will display dynamic instabilities, a phenomenon which is completely unaccounted for in any existing treatment of manipulation.

We attributed this inconsistency, i.e., the fact that a learned-then-manipulated causal model is not equal to the manipulated-then-learned model, to an inappropriate use of the Manipulation Postulate in manipulation formalisms. In explaining the inconsistency, we applied the work of Iwasaki and Simon (1994), which deals with representing causality in time-dependent systems based on structural equation models combined with differential equation systems. They show that physical systems possessing stable fixed points may possess multiple causal graphs depending on the time-scale being modeled. We show that the Manipulation Postulate applied to Iwasaki-Simon-type graphs for our two paradoxical systems, *modeled on an infinitesimal time-scale* (graphs which we refer to as "differential causal graphs"), produce equilibrium causal graphs with the correct independence relations. Furthermore, we show how these differential causal models correctly predict the presence of instabilities under manipulations of the system. We conclude that the Manipulation Postulate, and thus all existing manipulation formalisms, are only guaranteed to be valid on differential causal models.

Our result, perceived as rather controversial by the reviewers, is still unpublished. Our draft has been rejected three times by the Annual Conference on Uncertainty in Artificial Intelligence, which is a prestigious conference with a strict review process. Unfortunately, the conference format does not allow for resolving disagreements with the reviewers. We have been meeting and corresponding about our work with the leading experts in the field: Judea Pearl, Clark Glymour, Peter Spirtes, Greg Cooper and Herb Simon, who (similarly to the reviewers of the UAI conference) have been unable to demonstrate any major flaw in the paper. We believe that we are right and we are working on a submission to the *Journal of Artificial Intelligence Research*. A most recent draft of our paper is available from us (Dash & Druzdzel, 2000).

Interactive construction of causal graphical models based on causal mechanisms

Quality of decisions based on the decision-theoretic approach depends on the quality of the underlying models. Construction of these models is outside of the realm of both probability theory and decision theory and is usually very laborious. Aiding model building in computer systems can significantly reduce the model construction time while increasing model quality and can contribute to a wider applicability of decision theory in decision support systems.

We proposed an interactive approach to computer-aided model construction that builds on the concept of causal mechanisms. Causal mechanisms, which are local interactions among domain variables, are building blocks that determine the causal structure of a model. They are usually fairly well understood and model independent, and hence can be reused in different models. As they encode our understanding of local interactions and are fairly model independent, they can be easily reused in various models. When the algebraic form of the interaction is known, causal mechanisms are captured by so called structural equations. A model composed of causal mechanisms is causal and intuitive for human users. It also supports predictions of the effect of external interventions (decisions). As shown by Druzdzel and Simon (1993), conditional probability tables can be also viewed in causal models as descriptions of causal mechanisms. We assist users by identifying a set of mechanisms related to current model and let them choose from among them. In our knowledge-base, we encode mathematical relationships among the variables and, wherever known, the direction of causal influence among the variables. The mechanism-based view of model building is unique in the sense that it assists in building models that contain reversible causal mechanisms, i.e., mechanisms that work in several directions, depending on which of their variables are being manipulated at any given point. Building causal models is important for two reasons. Firstly, causal models are intuitive for human users to understand. Secondly, they allow for predicting the effect of external interventions, such as decisions.

We published the results of this work first in a 1998 Stanford AAAI Spring Symposium (Druzdzel, Lu & Leong, 1998) and then in the Annual 2000 Uncertainty in Artificial Intelligence conference (Li, Druzdzel & Leong, 2000).

Learning graphical models from data

Methods for learning probabilistic graphical models can be partitioned into at least two general classes: constraint-based search and Bayesian methods. The constraint-based approaches (Spirtes *et al.*, 1993, Verma & Pearl, 1991) search the data for conditional independence relations from which it is in principle possible to deduce the Markov equivalence class of the underlying causal graph. Two notable constraint-based algorithms are the PC algorithm which assumes that no hidden variables are present and the FCI algorithm which is capable of learning something about the causal relationships even assuming there are latent variables present in the data (Spirtes *et al.*, 1993). Bayesian methods (Cooper & Herskovits, 1991) utilize a search-and-score procedure to search the space of DAGs, and use the posterior density as a scoring function. There are many variations on Bayesian methods, however, most research has focused on the application of greedy heuristics, combined with techniques to avoid local maxima in the posterior density (e.g., greedy search with random restarts or best-first searches).

Both constraint-based and Bayesian approaches have advantages and disadvantages. Constraint-based approaches are relatively quick and possess the ability to deal with latent variables. However, constraint-based approaches rely on an arbitrary significance level to decide independencies, and they can be unstable in the sense that an error early on in the search can have a cascading effect that causes a drastically different graph to result. Bayesian methods can be applied even with very little data where conditional independence tests are likely to break down. Both approaches have the ability to incorporate background knowledge in the form of temporal ordering, or forbidden or forced arcs, but Bayesian approaches have the added advantage of being able to flexibly incorporate users' background knowledge in the form of prior probabilities over the structures and over the parameters of the network. In addition, Bayesian approaches are capable of dealing with incomplete records in the database. The most serious drawback to the Bayesian approaches is the fact that they are relatively slow.

Typically, Bayesian search procedures operate on the space of directed acyclic graphs (DAGs). However, recently researchers have investigated performing greedy Bayesian searches on the space of equivalence classes of DAGs (Spirtes, 1997, Madigan 1995, Chickering, 1996). The graphical objects representing equivalence classes have been called by several names ("patterns," "completed pdag representations," "maximally oriented graphs," and "essential graphs"). We use the term "essential graph" because we feel it is both descriptive and concise (but we acknowledge that the term "pattern" is more prevalent). An essential graph is a special case of a chain graph, possessing both directed and non-directed arcs, but no directed cycles. In order to specify an equivalence class it is necessary and sufficient to specify both a set of undirected adjacencies and a set of v-structures (a.k.a. "non-shielded colliders", a structure such as $X \rightarrow Y \leftarrow Z$ such that X is not adjacent to Z) possessed by the dag (Chickering, 1995). An essential graph therefore possesses undirected adjacencies when two nodes are adjacent, and it may possess directed adjacencies if a triple of nodes possesses a v-structure or if an arc is required to be directed due to other v-structures (Anderson, 1995). The space of essential graphs is smaller than the space of DAGs; therefore it is hoped that performing a search directly within this space might be beneficial; however, the Bayesian metric must be applied to a DAG, therefore these procedures incur the additional cost required to convert back and forth between essential-graph-space and DAG-space. Results from the above work have shown to be promising, however.

Researchers have also developed two-stage hybrid algorithms, where the first stage performs a constraint-based search and uses the resulting graph as input into a second-stage Bayesian search. In particular, (Singh, 1993) used

the PC algorithm to generate an absolute temporal ordering on the nodes for use with the K2 algorithm (Cooper & Herskovits, 1992), which requires such an ordering on the input (Spirtes, 1997) use the PC algorithm to generate a good starting graph for use in their greedy search over the space of essential graphs.

Our insight into learning graphical models from data led us to the development of a hybrid constraint-based/Bayesian algorithm for learning causal networks in the presence of sparse data. The algorithm searches the space of equivalence classes of models (essential graphs) using a heuristic based on conventional constraint-based techniques. Each essential graph is then converted into a directed acyclic graph and scored using a Bayesian scoring metric. Two variants of the algorithm are developed and tested using data from randomly generated networks of sizes from 15 to 45 nodes with data sizes ranging from 250 to 2000 records. Both variations are compared to, and found to consistently outperform two variations of greedy search with restarts. This algorithm was presented in the 1999 Annual Conference on Uncertainty in Artificial Intelligence (Dash & Druzdzel, 1999).

Other contributions

Relevance-based methods in algorithms for Bayesian networks

Relevance reasoning in Bayesian networks can be used to improve efficiency of belief updating algorithms by identifying and pruning those parts of a network that are irrelevant for the computation. Relevance reasoning is based on the graphical property of d-separation and other simple and efficient techniques, the computational complexity of which is usually negligible when compared to the complexity of belief updating in general.

We used relevance reasoning in a belief updating algorithm for Bayesian networks that is applicable in practical systems in which observations are interleaved with belief updating. Our technique invalidates the posterior beliefs of those nodes that depend probabilistically on the new evidence and focuses the subsequent belief updating on the invalidated beliefs rather than on all beliefs. Very often observations invalidate only a small fraction of the beliefs and our scheme can then lead to substantial savings in computation. We reported the results of this work in 1998 FLAIRS conference (Lin & Druzdzel, 1998) and in the *International Journal of Pattern Recognition and Artificial Intelligence* (Lin & Druzdzel, 1999).

Hepar II medical diagnostic system

In order to demonstrate the usefulness of our system in practical setting, we have started a successful collaboration focusing on building a practical medical system for diagnosis of liver disorders. The resulting system, Hepar II uses our software at its core and consists of a Bayesian network model comprising over 60 variables, such as disorder variables, risk factors for various disorders, symptoms, and test results (Figure 5 shows the model). The system's parameters are obtained from a database of real patient cases collected at the Institute of Food and Feeding in Warsaw, Poland. The resulting system will be applied both as a diagnostic tool in clinical setting and as a tool for training beginning diagnosticians. The result of our work have been published in several conferences, workshops, and symposia (listed in the publication list). We are working on a submission of this paper to a medical informatics journal.

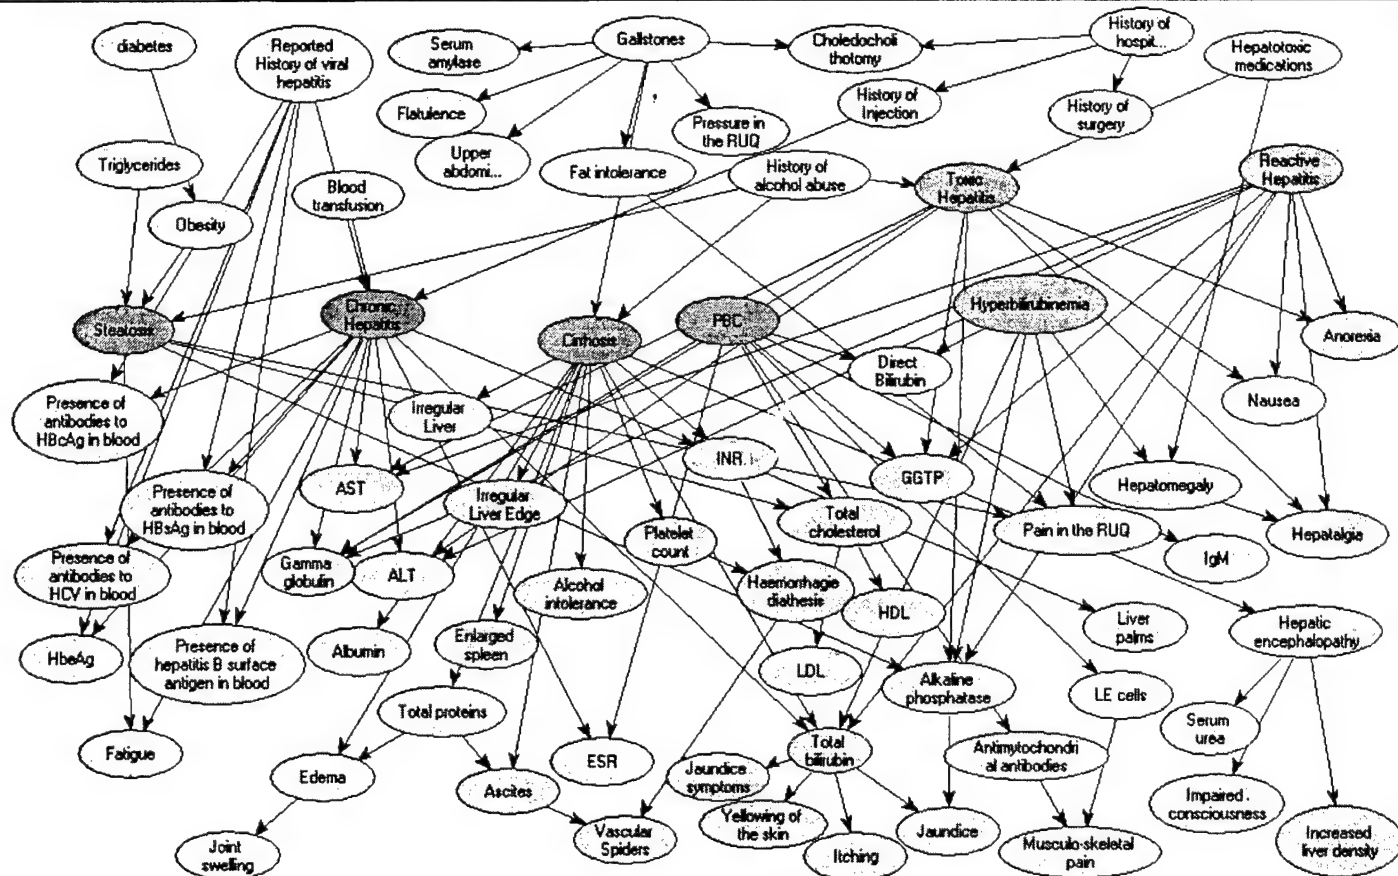


Figure 5: The Hepar II Bayesian network model.

GeNIe and SMILE[☺]

A major accomplishment that originates from the project is the implementation of the system. Since there is much interest now in Bayesian networks, influence diagrams, and decision-analytic systems, we have put much effort in making the implementation easy to use and robust and decided to share it with the community. We believe that this will bring a high payoff in the long run in terms of practical applications based on our system. We have written a comprehensive on-line help for **GeNIe** (the user interface running on Windows machines), useful for both beginning modelers and students in decision-analytic methods and a documentation for **SMILE[®]** (**S**tructural **M**odeling, **I**nference, and **L**earning **E**ngine), a portable library of C++ classes for decision-theoretic reasoning, **GeNIe**'s reasoning engine. We have also developed **SmileX**, an Active-X control version of **SMILE[®]** that allows the program to be used from most Windows applications, including Visual Basic, Java, Excel, and HTML pages. We have made our programs available on the World Wide Web in July 1998 (the address to download the program is: <http://www2.sis.pitt.edu/~genie>). Over 2,300 people from countries all over the world downloaded it since the release date. We have heard very positive feedback from these users. We have presented the programs in a number of research lectures and in conferences, including the American Association for Artificial Intelligence conference (Druzdzel, 1999a) and the American Medical Informatics Association (AMIA) conference (Druzdzel, 1999b). A screen shot of **GeNIe** is presented in Figure 6.

We have also implemented a module for assistance in model building based on causal mechanisms, a specialized module for diagnosis, and a module for learning models from data. These modules have not been released on the World Wide Web yet because they are not sufficiently reliable (given the large number of users of our programs, we have adopted high quality standards for releasing our software).

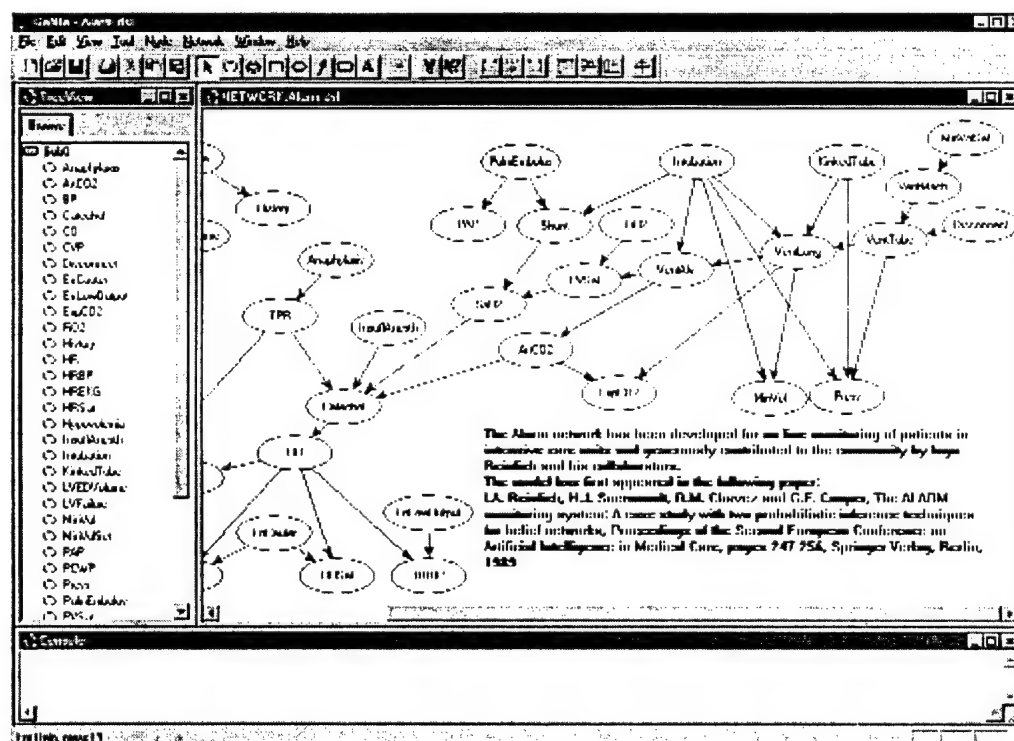


Figure 6: A screen shot of **GeNIe**.

Publications acknowledging support from AFOSR

Journals:

Jian Cheng and Marek J. Druzdzel. BN-AIS: An adaptive importance sampling algorithm for evidential reasoning in large Bayesian networks. *Journal of Artificial Intelligence Research* (to appear).

Marek J. Druzdzel and Hans van Leijen. Causal reversibility in Bayesian networks. *Journal of Experimental and Theoretical Artificial Intelligence* (to appear).

Marek J. Druzdzel and Linda C. van der Gaag. "Where Do the Numbers Come From?": Guest Editors' Introduction. Introduction to *IEEE Transactions on Knowledge and Data Engineering*, special issue on building probabilistic models (to appear).

Yan Lin and Marek J. Druzdzel. Relevance-based incremental belief updating in Bayesian networks. *International Journal of Pattern Recognition and Artificial Intelligence (IJPRAI)*, 13(2):285-295, 1999.

Major peer reviewed conferences:

Jian Cheng and Marek J. Druzdzel. Computational investigation of low-discrepancy sequences in Bayesian networks. In *Proceedings of the Sixteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-2000)*, pages 72-81, Morgan Kaufmann Publishers, Inc., San Francisco, CA, 2000.

Tsai-Ching Lu, Marek J. Druzdzel and Tze-Yun Leong. Causal mechanism-based model construction. In *Proceedings of the Sixteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-2000)*, pages 353-362, Morgan Kaufmann Publishers, Inc., San Francisco, CA, 2000.

Haiqin Wang and Marek J. Druzdzel. User interface tools for navigation in conditional probability tables and elicitation of probabilities in Bayesian networks. In *Proceedings of the Sixteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-2000)*, pages 617-625, Morgan Kaufmann Publishers, Inc., San Francisco, CA, 2000.

Jian Cheng and Marek J. Druzdzel. Latin hypercube sampling in Bayesian networks. In *Proceedings of the Thirteenth International Florida Artificial Intelligence Research Symposium (FLAIRS-2000)*, Special Track on Uncertain Reasoning.

Marek J. Druzdzel. **GeNIe**: A development environment for graphical decision-analytic models. In *Proceedings of the 1999 Annual Symposium of the American Medical Informatics Association (AMIA-1999)*, page 1206, Washington, D.C., November 6-10, 1999.

Marek J. Druzdzel, Agnieszka Onisko, Daniel Schwartz, John N. Dowling and Hanna Wasyluk. Knowledge engineering for very large decision-analytic medical models. In *Proceedings of the 1999 Annual Symposium of the American Medical Informatics Association (AMIA-1999)*, page 1049, Washington, D.C., November 6-10, 1999.

Marek J. Druzdzel. **SMILE**[®]: Structural Modeling, Inference, and Learning Engine and **GeNIe**: A Development environment for graphical decision-theoretic models (Intelligent Systems Demonstration). In *Proceedings of the Sixteenth National Conference on Artificial Intelligence (AAAI-99)*, pages 902-903, AAAI Press/The MIT Press, Menlo Park, CA, 1999.

Denver H. Dash and Marek J. Druzdzel. A hybrid anytime algorithm for the construction of causal models from sparse data. In *Proceedings of the Fifteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-99)*, pages 142-149, Morgan Kaufmann Publishers, Inc., San Francisco, CA, 1999.

Yan Lin and Marek J. Druzdzel. Relevance-based sequential evidence processing in Bayesian networks. In *Proceedings of the Uncertain Reasoning in Artificial Intelligence track of the Eleventh International Florida Artificial Intelligence Research Symposium Conference (FLAIRS-98)*, pages 446-450, AAAI Press/The MIT Press, Menlo Park, CA, 1998.

Other peer reviewed conferences, symposia, workshops, and book chapters:

Agnieszka Onisko, Marek J. Druzdzel and Hanna Wasyluk. Learning Bayesian network parameters from small data sets: Application of Noisy-OR gates. To appear in *Working Notes of the CaNew'2000 Workshop, European Conference on Artificial Intelligence*, Berlin, Germany, August 2000.

Marek J. Druzdzel and Roger R. Flynn. Decision Support Systems. To appear in Allen Kent (ed.) *Encyclopedia of Library and Information Science*, Marcel Dekker, Inc., 2000.

Marek J. Druzdzel and F. Javier Diez. Criteria for combining knowledge from different sources in probabilistic models. In *Working Notes of the workshop on "Fusion of Domain Knowledge with Data for Decision Support," Sixteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-2000)*, pages 23-29, Stanford, CA, 30 June 2000.

Agnieszka Onisko, Marek J. Druzdzel and Hanna Wasyluk. Extension of the Hepar II Model to Multiple-Disorder Diagnosis. In *Intelligent Information Systems*, M. Klopotek, M. Michalewicz, S.T. Wierzchon (eds.), pages 303-313, *Advances in Soft Computing Series*, Physica-Verlag (A Springer-Verlag Company), Heidelberg, 2000.

Agnieszka Onisko, Marek J. Druzdzel and Hanna Wasyluk. A Bayesian network model for diagnosis of liver disorders. In *Proceedings of the Eleventh Conference on Biocybernetics and Biomedical Engineering*, pages 842-846, Warsaw, Poland, December 2-4, 1999.

Marek J. Druzdzel and Clark Glymour. Causal inferences from databases: Why universities lose students. In Clark Glymour and Gregory F. Cooper (eds.), *Computation, Causation, and Discovery*, Chapter 19, pages 521-539, AAAI Press, Menlo Park, CA, 1999.

Denver H. Dash and Marek J. Druzdzel. A fundamental inconsistency between equilibrium causal discovery and causal reasoning formalisms. To appear in *Working Notes of the Workshop on Conditional Independence Structures and Graphical Models*, pages 17-18, Fields Institute, Toronto, Canada, 27 September - 1 October, 1999.

Marek J. Druzdzel. ESP: A mixed initiative decision-theoretic decision modeling system. In *Working Notes of the AAAI-99 Workshop on Mixed-initiative Intelligence*, pages 99-106, Orlando, Florida, 18 July 1999.

Yan Lin and Marek J. Druzdzel. Stochastic sampling and search in belief updating algorithms for very large Bayesian networks. In *Working notes of the AAAI-1999 Spring Symposium on Search Techniques for Problem Solving Under Uncertainty and Incomplete Information*, pages 77-82, Stanford, CA, March 22-24, 1999.

Agnieszka Onisko, Marek J. Druzdzel and Hanna Wasyluk. Graphical probabilistic models in diagnosis of liver disorders. In *Working Notes of the Third International Seminar on Statistics and Clinical Practice (45th Seminar of the International Centre of Biocybernetics)*, Warsaw, Poland, June 24-27, 1998.

Agnieszka Onisko, Marek J. Druzdzel and Hanna Wasyluk. A probabilistic causal model for diagnosis of liver disorders. In *Proceedings of the Seventh Symposium on Intelligent Information Systems (IIS-98)*, pages 379-387, Malbork, Poland, June 15-19, 1998.

Marek J. Druzdzel, Tsai-Ching Lu and Tze-Yun Leong. Interactive construction of decision models based on causal mechanisms. In *Working notes of the AAAI 1998 Spring Symposium on Interactive and Mixed-initiative Decision-theoretic Systems*, pages 38-44, Stanford, CA, March 23-25, 1998.

Hans van Leijen and Marek J. Druzdzel. Reversible causal mechanisms in Bayesian networks. In *Working notes of the AAAI 1998 Spring Symposium on Prospects for a Commonsense Theory of Causation*, pages 24-30, Stanford, CA, March 23-25, 1998.

Agnieszka Onisko, Marek J. Druzdzel and Hanna Wasyluk. Application of Bayesian belief networks to diagnosis of liver disorders. In *Proceedings of the Third Conference on Neural Networks and Their Applications*, pages 730-736, Kule, Poland, October 14-18, 1997.

Other papers:

Marek J. Druzdzel, Agnieszka Onisko, Daniel Schwartz, John N. Dowling and Hanna Wasyluk. Knowledge engineering for very large decision-analytic medical models. *Research Report CBMI-99-26, Center for Biomedical Informatics*, University of Pittsburgh, September 1999 (a full version of the short paper published in AMIA-99).

Agnieszka Onisko, Marek J. Druzdzel and Hanna Wasyluk. A Bayesian network model for diagnosis of liver disorders. *Research Report CBMI-99-27, Center for Biomedical Informatics*, University of Pittsburgh, September 1999.

Interactions / Transitions

Here are some of the applications of our results and our software:

The Decision Support Department of the United States Naval War College, Newport, RI, is using **GeNIe** and **SMILE**® in supporting a joint US NWC/US NAVEUR project on detection of sources of regional instabilities. The point of contact there is Bradd C. Hayes (hayesb@nwc.navy.mil).

Rockwell International Science Center, Palo Alto Laboratory, in collaboration with US Air Force Rome Laboratories are applying **GeNIe**, **SMILE**® and **SmileX** to the problem of battle damage assessment. The contact persons there are Mark Peot (peot@rpal.rockwell.com) and John F. Lemmer (John.Lemmer@rl.af.mil).

GeNIe and **SMILE**® have been applied in an intelligent tutoring system for teaching elementary physics, developed at University of Pittsburgh's Learning Research and Development Center (contact person is Prof. Kurt van Lehn, vanlehn@cs.pitt.edu). The system will be applied in teaching Navy cadets.

Dr. Wojtek Przytula (wojtek@hrl.com) at the Hughes Raytheon Laboratories uses **GeNIe** and **SMILE**® in a diagnostic system for General Motors Diesel locomotives.

We have two current points of contact who are interested in using the results of our work when our system implements both Bayesian networks and structural equations: Dr. Patrick Love at the Aluminum Company of America (ALCOA) Technical Center (Patrick.Love@alcoa.com), for strategic business planning at Aluminum Company of America, and Mr. Jeffrey Bolton (jb5c+@andrew.cmu.edu) and Mr. Kevin Lamb (kl3g+@andrew.cmu.edu) at the Carnegie Mellon University's Office of Planning and Budget, for strategic planning of university operations. These contacts will be followed up when **GeNIe** and **SMILE**® implement both equations and Bayesian networks.

Honors / Awards

2000 Robert R. Korfhage award (with Jian Cheng), awarded school-wide for the best paper co-authored between a student and a faculty member.

1999 Robert R. Korfhage award (with Jian Cheng), awarded school-wide for the best paper co-authored between a student and a faculty member.